

AN EFFICIENT FPGA IMPLEMENTATION OF A FLEXIBLE JPEG2000 DECODER FOR DIGITAL CINEMA

Antonin Descampe^{1,a}, François Devaux¹, Gaël Rouvroy², Benoît Macq¹, and Jean-Didier Legat²

¹Communications and Remote Sensing Lab., Stevin Building

E-mail : {descampe, devaux, macq}@tele.ucl.ac.be

²Microelectronics Lab., Maxwell Building

E-mail : {rouvroy, legat}@dice.ucl.ac.be

Université catholique de Louvain

Place du Levant, B-1348 Louvain-la-Neuve
Belgium

ABSTRACT

The image compression standard JPEG2000 proposes a large set of features, useful for today's multimedia applications. Unfortunately, its complexity is greater than older standards. A hardware implementation brings a solution to this complexity for real-time applications, such as Digital Cinema. In this paper, a decoding scheme is proposed with two main characteristics. First, the complete scheme takes place in an FPGA without accessing any external memory, allowing integration in a secured system. Secondly, a customizable level of parallelization allows to satisfy a broad range of constraints, depending on the signal resolution.

1. INTRODUCTION

Development and diversification of computer networks as well as emergence of new imaging applications have highlighted various shortcomings in actual image compression standards, such as JPEG. The lack of resolution or quality scalability is clearly one of the most significant drawbacks. The new image compression standard JPEG2000 [1] enables such scalability : according to the available bandwidth, computing power and memory resources, different resolution and quality levels can be extracted from a single bit-stream. In addition to this, the JPEG2000 baseline (part I of the standard) also proposes other important features: good compression efficiency, even at very low bit rates, lossless and lossy compression using the same coder, random access to the compressed bit-stream, error resilience, region-of-interest coding. A comprehensive comparison of the norm with other standards, performed in [2], shows that from a functionality point of view JPEG2000 is a true improvement.

The techniques enabling all these features are a wavelet transform (DWT) followed by an arithmetic coding of each subband. The drawback of these techniques is that they are computationally intensive, much more for example than a cosine transform (DCT) followed by an Huffman coding, which are those used in JPEG [2]. This complexity can be a problem for real-time applications.

Digital Cinema is one of these real-time applications. As explained in [3], edition, storage or distribution of video data can largely take advantage of the JPEG2000 feature set. Moreover, a video format named Motion JP2000 has been designed, which encapsulates JPEG2000 frames and enables synchronization with audio data [4]. Nevertheless, a high output rate is required at the decoding process and in order to meet this real-time constraint, a dedicated implementation of the most complex parts of the algorithm is needed.

In this paper, a complete JPEG2000 decoder architecture intended for video decoding is proposed. It has been implemented in VHDL and synthesized in an FPGA (Xilinx XC2V6000 [5]). It

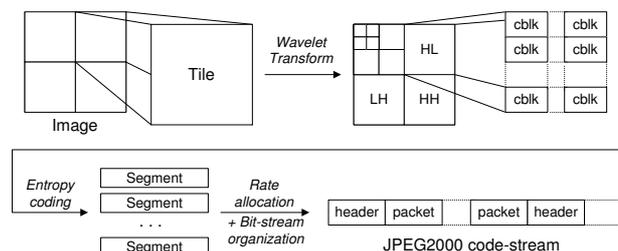


Figure 1: Coding steps of the JPEG2000 algorithm.

takes about 90% of the chip and the estimated frequency of operation is 90 Mhz. The proposed architecture decodes images line by line without accessing any external memory. It is highly parallelized and depending on available hardware resources, it can easily be adapted to satisfy various formats, from Digital Cinema to Video-on-Demand, and specific constraints like secure decoding, lossless capabilities, and higher precision (over 8 bits per pixel).

The rest of the paper is organized as follows. Section II briefly describes the JPEG2000 algorithm. In Section III, we present our decoder architecture as well as our implementation choices. The performance of the system is discussed in Section IV and the paper is concluded in Section V.

2. JPEG2000 OVERVIEW

In this Section, concepts and vocabulary useful for the understanding of the rest of the paper are presented. For more details, please refer to [1] or to [6]. Although a decoder architecture was achieved, *encoding* steps are explained here because their succession is easier to understand. Decoding process is achieved by performing these steps in the reverse order. Figure 1 presents the coding blocks which are explained below.

First of all, the image is split into rectangular blocks called tiles. They will be compressed independently from each other. An intra-components decorrelation is then performed on the tile: on each component a *discrete wavelet transform* is carried out. Successive dyadic decompositions are applied. Each of these uses a bi-orthogonal filter bank and splits high and low frequencies in the horizontal and vertical directions into four subbands. The subband corresponding to the low frequencies in the two directions (containing most of the image information) is used as a starting point for the next decomposition, as shown in Fig. 1. Two filter banks may be used : either the *Le Gall* (5,3) filter bank prescribed for lossless encoding or either the *Daubechies* (9,7) filter bank for lossy encoding.

Every subband is then split into rectangular entities called code-blocks. Each code-block will be compressed independently from

^a A.Descampe is funded by the Belgian NSF.

the others using a *context-based adaptative entropy coder*. It reduces the amount of data without losing information by removing redundancy present in the binary sequence. “Entropy” means it achieves this redundancy reduction using the probability estimates of the symbols. Adaptivity is provided by dynamically updating these probability estimates during the coding process. And “context-based” means the probability estimate of a symbol depends on its neighborhood (its “context”). Practically, entropy coding consists of

- *Context Modeling* : the code-block data is arranged in order to first encode the bits which contribute to the largest distortion reduction for the smallest increase in file size. In JPEG2000, the Embedded Block Coding with Optimized Truncation (EBCOT) algorithm [8] has been adopted to implement this operation. The coefficients in the code-block are bit-plane encoded, starting with the most significant bit-plane. Instead of encoding the entire bit-plane in one coding pass, each bit-plane is encoded in three passes with the provision of truncating the bit-stream at the end of each coding pass. During a pass, the modeler successively sends each bit that needs to be encoded in this pass to the Arithmetic Coding Unit described below, together with its context.
- *Arithmetic Coding* : the modeling step outputs are entropy coded using a MQ-coder, which is a derivative of the Q-coder. According to the provided context, the coder chooses a probability for the bit to encode, among predetermined probability values supplied by the JPEG2000 Standard and stored in a look-up table. Using this probability, it encodes the bit and progressively generates code-words, called segments.

During the *rate allocation* and *bit-stream organization* steps, segments from each code-block are scanned in order to find optimal truncation points to achieve various targeted bit-rates. Quality layers are then created using the incremental contributions from each code-block. Compressed data corresponding to the same component, resolution, spatial region and quality layer is then inserted in a packet. Packets, along with additional headers, form the final JPEG2000 code-stream.

3. PROPOSED ARCHITECTURE

In this section, we first present the constraints we used for our JPEG2000 decoder architecture. Implementation choices made in order to meet these constraints are then explained. Finally, the complete architecture is presented.

3.1 Constraints

As our decoder is designed for real-time video processing, three main constraints have been identified :

- *High output bit-rate* : all implementation choices have been made in order to increase this bit-rate. With the Xilinx XC2V6000 used, we wanted our architecture to satisfy at least the 1080/24p HDTV format. This means an output rate of about 1200 megabits per second (Mbps) for 8-bit 4:4:4 images.
- *Security* : no data flow may transit outside of the FPGA if it is not crypted or watermarked. This constraint enables a completely secured decoding scheme, as the decompression block might be inserted between a decryption block and a watermark block, all these three blocks being in the same FPGA (Fig. 2).
- *Flexibility* : computationally intensive parts of the decoding process must be independent blocks which can easily be duplicated and parallelized. This allows the proposed architecture to satisfy a broad range of output bit-rate constraints and therefore to be easily adapted to upcoming Digital Cinema standards.

3.2 Implementation choices

To meet these constraints, the following implementation choices have been made.

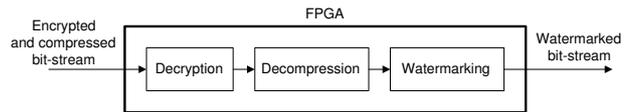


Figure 2: A secured decoding scheme.

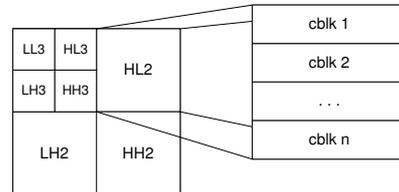


Figure 3: Customized code-block dimensions.

No external memory has been used which meets the security constraint and also increases the output bit-rate. As internal memory resources are limited, large image portions cannot be stored and the decoding process must be achieved in a line-based mode.

In order to increase the output bit-rate, three *parallelization* levels have been used. The first one is a duplication of the entire architecture which allows various tiles to be decoded simultaneously. The second parallelization level tries to compensate the compute load difference between the entropy decoding unit (EDU) and the inverse wavelet transform (IDWT). The EDU is indeed much more complex than the IDWT and must therefore be parallelized. This is possible as each code-block is decoded independently from the others. Finally, a third level of parallelization, known in the JPEG2000 standard as the parallel mode, is obtained inside each EDU. By default, each bit-plane is decoded in three successive passes but specifying some options ([7], p.508) during the encoding process makes it possible to decode the three passes simultaneously. This implies that each EDU contains one Context Modeling Unit (CMU) and three Arithmetic Decoding Units (ADU).

Another option specified during the encoding process that increases the output bit-rate of the decoder is the *bypass mode* ([7],p.504). The more correlated the probability estimates of the bits to encode are, the more efficient the ADU is. This is especially the case in the most significant bit-planes while the last bit-planes are most of the time totally uncorrelated. With the bypass mode enabled, these last bit-planes are therefore raw-coded¹.

Some choices about *image partitioning* have also been made. A 512x512 tile partition avoids an external memory use and enables one of the parallelization level mentioned above. Inside each tile, even if the code-block maximum size specified in the norm is 4096 pixels, code-blocks in our implementation do not exceed 2048 pixels. As we shall see, this implies no significant efficiency loss but allows a 50% memory resources saving. Furthermore, the code-block dimensions have been chosen so that each of them systematically covers the width of the subband to which it belongs (Fig. 3). As the IDWT processes the subband data line by line, such code-block dimensions enables a line-based approach of the overall process, reducing the size of the image portions to store between EDU and IDWT.

These last implementation choices (parallel mode, bypass mode and image partitioning) imply an efficiency loss during the encoding process. Table 1 shows the corresponding *psnr* losses for various compression ratio. In comparison to the improvements provided by these choices, quality losses are quite reduced, especially for small ratios which are the ones used in the targeted applications.

Another choice has to be made in order to enable a line-based

¹This means they are inserted “as is” in the bit-stream.

Compression ratio	PSNR [dB]	
	Default options	Options used
1:50	37,36	35,83 (-1,53)
1:25	40,10	38,74 (-1,36)
1:10	43,25	42,40 (-0,85)
1:5	45,85	45,31 (-0,54)

Table 1: Average PSNR for a set of images 1920x1080, 8 bpp

processing of the image. To reconstruct one line of the original image, the IDWT needs the corresponding line at each resolution. In order to minimize the image portions size to store, data inside the bit-stream is organized so that the whole compressed data corresponding to a specific spatial region of the image is contiguous in the bit-stream. Various progression orders are allowed during the JPEG2000 encoding process and one of them enables such kind of feature.

A last implementation choice aims at achieving some lightweight operations in software. These operations are indeed essentially data handling and are easily implemented using pointers in C. To keep the decoding process secure, headers and markers (needed by these operations) are not crypted and only the packet bodies are.

As it can be seen, some options, known by any universal JPEG2000 encoder, must be specified during the encoding process. Our architecture is unable to decode a JPEG2000 code-stream that has not been encoded using these options. As this architecture is dedicated to decode video data at the highest output bit-rate, we did not consider it efficient to realize a universal decoder.

3.3 Architecture

Figure 4 presents the hardware part of our architecture. Each EDU contains three ADU's which reflects the parallel mode. The bypass mode is also illustrated by the bypass line under each ADU. The Dispatch IN and OUT blocks are used to dissociate the entropy decoding step from the rest of the architecture and enable the flexibility mentioned above.

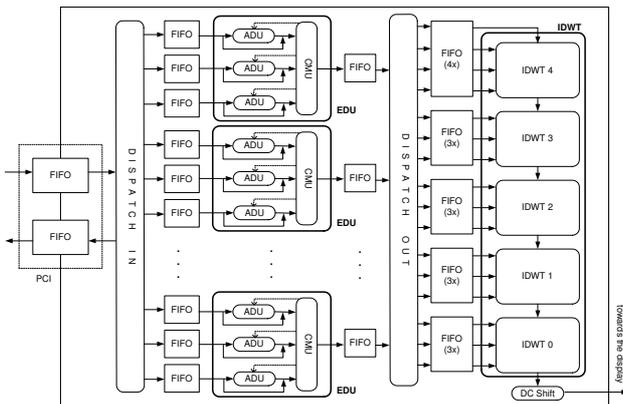


Figure 4: Proposed architecture.

When Dispatch IN receives a new JPEG2000 code-stream from the PCI, it chooses one of the free EDU's and connects the data stream to it. Dispatch OUT retrieves decompressed data from each EDU and connects it to the correct subband FIFO. In this way, a maximum of EDU's is always used simultaneously.

Instead of explaining the whole architecture in details, we present below some characteristics from the ADU, CMU and IDWT blocks that are worth being noticed.

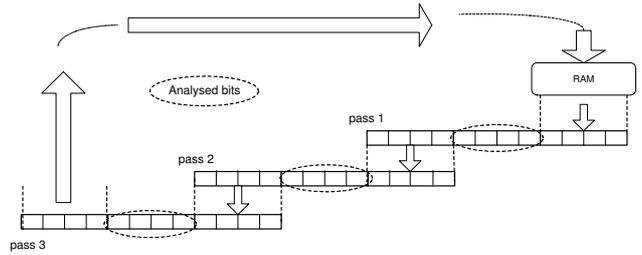


Figure 5: State variables transfer between the three passes.

3.3.1 Context Modeling Unit

During the Context Modeling step, each code-block is decoded along its bit-planes. As mentioned above, the CMU decodes the bits of a given bit-plane in three passes, each bit being decoded by one of these three passes. Once the CMU decides to decode a bit, it sends its context to the ADU and waits for the decoded bit. In order to determine in which pass a bit need to be decoded, various state variables are held up to date by the CMU.

The CMU architecture is based on the one developed by Andra *et al.* in [9] and proposes various optimizations, most of them due to the use of the parallel mode.

- As shown in Fig. 5, several registers are used to communicate the state variables from one pass to another, offering significant memory reduction.
- A unique counter for the three passes manages the code-block's characteristics, which enables highly simplified control parts for the three blocks. It synchronizes the passes and controls the code-block's borders. This allows the three entities to work the same way either they decode bits inside a code-block or either on its border.
- As the EBCOT algorithm complexity makes the CMU the slowest component of the decoder, the memories designed to receive the decoded code-blocks are able to take care of two code-blocks at a time. In this way, the EDU can begin to decode a new code-block while the previous one is still being processed by the IDWT.

3.3.2 Arithmetic Decoding Unit

During the Arithmetic Decoding step, the ADU uses the context provided by the CMU, together with the compressed bit-stream, to output the decoded bit. The proposed architecture has several original characteristics.

- The main control state machine consists of only 5 states. Furthermore, the CMU is waiting for the ADU answer during only three of them. Therefore a symbol may be decoded in 3 clock cycles. The bypass mode still improves this result.
- In a direct implementation of the JPEG2000 Standard [1], the bit-stream enters the ADU one bit at a time (left-shift operations). In our architecture, a speculative computation of the number of left-shifts is performed and all the shifts are then realized in one clock cycle.
- The compressed data loading is performed in an independent process and during non-critical moments of the whole decoding process, i.e. when the CMU is not waiting for an answer.

3.3.3 Inverse DWT

In JPEG2000, the DWT is implemented using a lifting-based scheme [10]. Compared to a classic implementation, it reduces the computational and memory cost, allowing in-place computation.

To reconstruct one resolution level, an horizontal transformation is applied first, followed by a vertical one. In our architecture, further detailed in [11], both transformations are entirely pipeline : every new coefficient "pushes" the already present ones a little more

Slices	30,323 over 33,792 (89.7%)
Look-Up Tables	51,416 over 67,584 (76.1%)
RAM blocks (16kbits)	89 over 144 (61.8%)
CLK1 (EDU's & Dispatch)	89.9 MHz
CLK2 (IDWT)	75.9 MHz

Table 2: Synthesis results of the decoding scheme in a Xilinx XC2V6000

Compression ratio	10-EDU [Mbps]	IDWT [Mbps]	Complete Scheme [#imgs(1920x1080)/sec]
1:10	728.0	2 440	14.63
1:20	1 290	2 440	25.92
1:32	2 137	2 440	42.94

Table 3: Bit-rates achieved by the Proposed architecture

through the pipeline, toward the output. Moreover, an efficient implementation of the lifting scheme allows to buffer only two entire lines of the level being reconstructed, enabling a line-based image processing.

The whole IDWT architecture has already been presented in Fig. 4. In comparison with Chrysafis who presented in [12] such kind of architecture, various optimizations have been made. First, as mentioned above, the lifting scheme has been adopted for each level. Second, interconnection of the blocks has been carefully studied and simplified. Each IDWT i -block ($i = 0..4$) reconstructs one resolution level and behaves at its output like a FIFO. Therefore, each block “sees” four FIFO’s as its inputs. Finally, the pipeline characteristic, already present inside each level has been extended to the whole architecture. Thanks to the progression order chosen, the sixteen FIFO’s (one per subband) are filled as uniformly as possible. As soon as its input FIFO’s contains data (including the one simulated by the preceding level), an IDWT i -block begins to reconstruct its level. When the pipeline is full, coefficients of the reconstructed image are provided line by line at each clock cycle.

4. PERFORMANCES

The architecture presented has been implemented in VHDL and synthesized and routed in an FPGA (Xilinx XC2V6000) using 10 EDU in parallel. Table 2 presents the resources used with this configuration.

Table 3 presents the bit-rates achieved by our architecture with the configuration described above. 24bpp-images were encoded using options explained in section 3.2. As we can see, this configuration yet enables real-time 8-bit 4:4:4 video decoding for the 1080/24p HDTV format and a compression ratio of 20. For a compression ratio of 11, the same format is supported with 8-bit 4:2:2 images.

Several other JPEG2000 hardware implementations have been developed. The main differences between two recent ones and the proposed architecture are listed in Table 4.

	Barco Silex[13]	Arizona Univ.[14]	Proposed architecture
Technology	FPGA XC2V3000	ASIC 0.18 μ m	FPGA XC2V6000
Tile size	128x128	128x128	512x512
Cblk size (max.)	32x32	32x32	64x32
Wavelet filters used	(5,3)-lossless (9,7)-lossy	(5,3)-lossless (9,7)-lossy	(5,3) lossy and lossless
Entropy coders	8	3	10

Table 4: Differences between two recent implementations and the proposed architecture

5. CONCLUSION

In this paper, we have proposed a hardware JPEG2000 decoder for real-time applications such as Digital Cinema. It has been implemented in VHDL, and synthesized and routed in an FPGA.

Various previous contributions have been joined together and optimized to provide a complete, flexible, secure, high performance decoding scheme.

The system proposed is secure because no external memory is used and the data flow is protected during the whole decoding process.

Thanks to three different levels of parallelization and a line-based data processing, high output rates are achieved. With a compression ratio of 20, the configuration synthesized in the FPGA supports the 1080/24p HDTV format for 8-bit 4:4:4 images.

Finally, the system proposed is highly flexible. In order to satisfy a broad range of constraints, including upcoming standards, two of the three parallelization levels are very easily customizable. They allow the proposed architecture to fit in any FPGA without further development.

REFERENCES

- [1] *ISO/IEC 15444-1: Information Technology-JPEG 2000 image coding system-Part 1: Core coding system*, 2000.
- [2] D. Santa-Cruz, R. Grosbois, and T. Ebrahimi, “JPEG2000 performance evaluation and assessment”, *Signal Processing: Image Communication*, vol. 17, no. 1, pp. 113-130, January 2002.
- [3] S. Foessel, “Motion JPEG2000 and Digital Cinema”, ISO/IEC JTC 1/SC 29/WG1 N2999, July 2003.
- [4] *ISO/IEC 15444-3: Information Technology-JPEG 2000 image coding system-Part 3: Motion JPEG 2000*, 2002.
- [5] VirtexTM-II platform FPGAs: Complete Data Sheet. Xilinx. [Online]. Available: <http://www.xilinx.com>.
- [6] M. Rabbani and R. Joshi, “An overview of the JPEG2000 still image compression standard”, *Signal Processing: Image Communication*, vol. 17, no. 1, pp. 3-48, January 2002.
- [7] D. Taubman and M. W. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards and Practice*, Kluwer Academic, Boston, MA, USA, 2002.
- [8] D. Taubman, “High performance scalable image compression with EBCOT”, *IEEE Trans. on Image Processing*, vol. 9, no. 7, pp. 1158-1170, July 2000.
- [9] K. Andra, T. Acharya, and C. Chakrabarti, “Efficient VLSI implementation of bit plane coder of JPEG2000”, in *Proc. SPIE Int. Conf. Applications of Digital Image Processing XXIV*, vol. 4472, pp. 246-257, December 2001.
- [10] I. Daubechies and W. Sweldens, “Factoring wavelet transforms into lifting steps”, *J. Fourier Anal. Applic.*, vol. 4, pp. 247-269, 1998.
- [11] G. Dillen and B. Georis, “JPEG 2000 : étude et conception du décodeur arithmétique et de la transformée en ondelettes”. Microelectronics Laboratory (DICE), UCL, Belgium, June 2001.
- [12] C. Chrysafys and A. Ortega, “Line based, reduced memory, wavelet image compression”, *IEEE Trans. on Image Processing*, vol. 9, no. 3, pp. 378-389, March 2000.
- [13] JPEG2000 Decoder: BA111JPEG2000D Factsheet. Barco-Silex, October 2003. [Online]. Available: <http://www.barco.com>.
- [14] K. Andra, T. Acharya, and C. Chakrabarti, “A High-Performance JPEG2000 Architecture”, *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 13, no. 3, pp. 209-218, March 2003.